

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of the claims in the application.

1. (Currently amended) A method, comprising:

obtaining a plurality of e-mails intended for distribution to a plurality of respective destinations; [[and]]

creating a data node for each e-mail in said plurality of e-mails, wherein each data node includes a pointer to the corresponding e-mail in persistent storage;

processing the plurality of e-mails data nodes solely within non persistent storage, without requiring that information indicative of the e-mails be written to and then read from persistent storage during the processing of the data nodes e-mails, wherein said processing comprises, for each respective data node:

(i) determining a destination domain of the respective data node;

(ii) adding the respective data node to a queue corresponding to the destination domain of the respective data node when the queue exists; and

(iii) creating a queue corresponding to the destination domain and adding the respective data node to the created queue when the queue does not exist; and wherein said processing further comprises:

selecting a first queue that contains data nodes;

retrieving e-mails corresponding to each of the data nodes in the first queue;

sending each of the retrieved e-mails corresponding to each of the data nodes in the first queue to a destination domain of the first queue; and

extinguishing the first queue.

2. (Previously presented) A method as in claim 1, further comprising

storing, in persistent storage, recovery information indicative of the processing, said recovery information being used for recovery from a system fault.

3. (Currently amended) A method as in claim 2, wherein said recovery information includes information indicative of [[a]] the plurality of e-mails, wherein said information is indicative of less than the entirety of each e-mail in said plurality of e-mails.

4. (Currently amended) A method as in claim 3, wherein said information indicative of an e-mail in the plurality of e-mails includes a bit vector.
5. (Currently amended) A method as in claim 1, wherein ~~said processing comprises:~~
arranging information about the plurality of e-mails into a plurality of queues, each queue representing a single domain; and
said sending of each of the e-mails corresponding to each of the data nodes in the first queue to a recipient, by sending a plurality of e-mails from ~~[[a]]~~ the first queue in said plurality of queues to the single destination domain is done ~~that the queue represents~~, at a specific sending instance.
6. (Currently amended) A method as in claim 5, wherein said sending comprises ~~[[:]]~~ opening a communication channel to a single specified domain ~~[[:]]~~ and sending each of the a plurality of e-mails within the single communication channel.
7. (Previously presented) A method as in claim 3, wherein said recovery information includes a numerical designation for each e-mail in said plurality of e-mails, and a state of processing of each e-mail in said plurality of e-mails.
8. (Cancelled)
9. (Currently amended) A method as in claim ~~[[8]]~~ 1, wherein said selecting comprises selecting a first queue which has the greatest number of the e-mails within the queue.
10. (Currently amended) A method as in claim ~~[[8]]~~ 1, wherein said selecting comprises selecting a first queue which has existed for the greatest period of time.
11. (Currently amended) A method as in claim ~~[[8]]~~ 1, further comprising, during said selection of said first queue selecting, asynchronously looking up single domain name server information for a second queue, ~~in said plurality of queues that is different than the first queue, and selecting the second queue.~~
12. (Currently amended) A method as in claim 1, ~~further comprising:~~

~~processing the plurality of e-mails by separating wherein the creating step~~
separates personalized information about each e-mail in the plurality of e-mails from non-personalized information.

13. (Previously presented) A method as in claim 12, wherein said non-personalized information includes e-mail destination information.

14. (Currently amended) A method as in claim 5, wherein said processing further comprises:

determining information about processing by said ~~single~~ destination domain; and
adjusting a speed of sending of the e-mails based on said information about processing of said ~~single~~ destination domain.

15. (Previously presented) A method as in claim 14, wherein said information about processing comprises a speed of e-mail processing.

16. (Currently amended) A method as in claim 1, further comprising:

maintaining a log representing information relating to a number of e-mails in said plurality of e-mails which have been processed; and

comparing contents of said log with licensing information, to determine if ~~said information~~ the number of ~~relating to~~ e-mails that has been processed exceeds a licensed number.

17. (Previously presented) A method as in claim 1, comprising:

storing recovery information about a state of processing of the plurality of e-mails to persistent storage, wherein said recovery information comprises less than the entirety of the plurality of e-mails; and

wherein the processing of the plurality of e-mails directs the plurality of e-mails to a desired location without writing the plurality of e-mails to persistent storage during said processing.

18. (Original) A method as in claim 17, wherein said processing comprises sending e-mails from an e-mail client to a desired location.

19. (Previously presented) A method as in claim 17, wherein said processing comprises receiving e-mails from and distributing said e-mails to said desired location.

20. (Previously presented) A method as in claim 17, wherein said recovery information includes information indicative of said plurality of e-mails, wherein said information is indicative of less than the entirety of each e-mail in said plurality of e-mails.

21. (Currently amended) A method as in claim 19, wherein said information indicative of ~~the~~ an e-mail in the plurality of e-mails includes a bit vector formed from ~~[[an]]~~ the e-mail, in said plurality of e-mails, that is indicative of the e-mail.

22. (Previously presented) A method as in claim 17, wherein said processing comprises:
arranging information about the e-mails into a plurality of queues, each queue in said plurality of queues representing a single domain; and
sending e-mails to a recipient, by sending a plurality of e-mails to a single domain, represented by a queue in said plurality of queues, at a specific sending instance.

23. (Currently amended) A method as in claim 18, wherein said sending comprises:
opening a communication channel to said desired location; and
sending a plurality of e-mails within the ~~single~~ communication channel.

24. (Previously presented) A method as in claim 17, wherein said recovery information includes a number of e-mails, and a state of processing of each e-mail in said number of e-mails.

25. (Previously presented) A method as in claim 22, further comprising
selecting a first queue in said plurality of queues to be processed, and
sending e-mails from the first queue all at once to the single domain represented by the first queue.

26. (Previously presented) A method as in claim 25, wherein said first queue has the most e-mails within the queue.

27. (Previously presented) A method as in claim 25, wherein said first queue has existed for the greatest period of time.

28. (Previously presented) A method as in claim 25, further comprising, during selection of said first queue,

asynchronously looking up domain name server information for a second queue in said plurality of queues that is different than the first queue.

29. (Previously presented) A method as in claim 17, further comprising:

processing the plurality of e-mails by separating personalized information about each e-mail in the plurality of e-mails from non-personalized information.

30. (Previously presented) A method as in claim 29, wherein said non-personalized information includes destination information for the plurality of e-mails.

31. (Previously presented) A method as in claim 22, wherein said processing comprises:

determining a speed of processing of said single domain; and

adjusting a speed of sending of the e-mails based on said speed of processing of said single domain.

32. (Previously presented) A method as in claim 17, further comprising:

maintaining a log representing information relating to e-mails which have been processed; and

comparing contents of said log with licensing information, to determine if said information relating to e-mails exceeds a licensed number.

33-59. (Cancelled)

60. (Currently amended) A method, comprising:

obtaining a plurality of e-mails for processing;

forming a queue map comprising a plurality of queues, each queue in the plurality of queues associated with a specific domain information about said plurality of e-mails, the queue map representing a plurality of destinations for the plurality of e-mails;

sending a plurality of e-mails to a specific destination in said plurality of destinations at a specific time; and

asynchronously looking up, during said sending step, DNS information for a domain name information using an asynchronous DNS resolver that operates from an offline DNS cache that is periodically updated, for a different destination in said plurality of destinations, to be sent at a future time.

61. (Previously presented) A method as in claim 60, further comprising:

processing the plurality of e-mails solely within non persistent storage, without requiring that information indicative of the plurality of e-mails be written to and then read from persistent storage during the processing of the plurality of e-mails.

62. (Previously presented) A method as in claim 61, further comprising:

storing, in persistent storage, recovery information indicative of the processing, wherein said recovery information is used for recovery from a system fault.

63. (Currently amended) A method as in claim 61, wherein said recovery information includes information indicative of a plurality of e-mails, wherein said information is indicative of less than the entirety of each of the ~~e-mail~~ e-mails in said plurality of e-mails.

64. (Currently amended) A method as in claim 60, wherein said processing comprises:

arranging information about the plurality of e-mails into [[a]] said plurality of queues, each queue in said plurality of queues representing a single domain; and

sending e-mails to a recipient, by sending a plurality of e-mails from a queue in said plurality of queues to the single domain that the queue represents at a specific sending instance.

65. (Previously presented) A method as in claim 64, wherein said sending comprises:

opening a communication channel to the single domain; and

sending a plurality of e-mails within the communication channel.

66. (Previously presented) A method as in claim 63, wherein said recovery information includes a number of e-mails, and a state of processing of each e-mail in said number of said e-mails.

67. (Previously presented) A method as in claim 64, further comprising:
selecting a first queue to be processed; and
sending e-mails from the first queue all at once to the single domain.
68. (Previously presented) A method as in claim 67, wherein said first queue has the most e-mails within the queue.
69. (Previously presented) A method as in claim 67, wherein said first queue has existed for the greatest period of time.
70. (Previously presented) A method as in claim 67, further comprising, during selection of said first queue, asynchronously looking up single domain name server information for a second queue that is different than the first queue.
71. (Previously presented) A method as in claim 64, wherein said sending further comprises:
determining a speed of processing of said domain; and
adjusting a speed of processing of the e-mails in the queue based on said speed of processing of said single domain.
72. (Previously presented) A method as in claim 60, further comprising:
maintaining a log representing a number of e-mails which have been sent; and
comparing contents of said log with licensing information, to determine if said number exceeds a licensed number.
73. (Currently amended) A method, comprising:
obtaining a plurality of e-mails for processing;
forming organization information about said plurality of e-mails, wherein said organization information represents a plurality of queues, each queue in said plurality of queues comprising e-mails in said plurality of e-mails that are intended for distribution to a common destination; and
selecting a first queue in said plurality of queues to send e-mails, based on characteristics of the e-mails in the first queue and, during the selecting step,

asynchronously looking up DNS information for a domain name using an asynchronous DNS resolver that operates from an offline DNS cache that is periodically updated, for a second queue in said plurality of queues, different than the first queue.

74. (Previously presented) A method as in claim 73, further comprising:

processing the plurality of e-mails solely within non persistent storage, without requiring that information indicative of the plurality of e-mails be written to and then read from persistent storage during processing.

75. (Previously presented) A method as in claim 73, wherein said first queue has the most e-mails within the queue.

76. (Previously presented) A method as in claim 73, wherein said first queue has existed for the greatest period of time.

77. (Cancelled)

78. (Original) A method as in claim 73, further comprising storing, in persistent storage, recovery information indicative of the processing, said recovery information being used for recovery from a system fault.

79. (Currently amended) A method as in claim 78, wherein said recovery information includes information indicative of [[a]] said plurality of e-mails, wherein said information is indicative of less than the entirety of each e-mail in said plurality of e-mails.

80. (Currently amended) A method as in claim 73, wherein said processing comprises:

arranging information about the plurality of e-mails into a plurality of queues, each queue in the plurality of queues representing a single domain; and

sending e-mails to a recipient, by sending a plurality of e-mails to a single domain at a specific sending instance.

81. (Previously presented) A method as in claim 80, wherein said sending comprises:

opening a communication channel to the single domain; and

sending a plurality of e-mails within the communication channel.

82. (Previously presented) A method as in claim 80, wherein said processing comprises:
determining a speed of processing of said single domain; and
adjusting a speed of sending of e-mails to said single domain based on said speed of processing of said single domain.

83. (Previously presented) A method as in claim 73, further comprising:
maintaining a log representing a number of e-mails which have been processed;
and
comparing contents of said log with licensing information, to determine if said number exceeds a licensed number.

84. (New) A computer system comprising:
means for obtaining a plurality of e-mails intended for distribution to a plurality of respective destinations;
means for creating a data node for each e-mail in said plurality of e-mails, wherein each data node includes a pointer to the corresponding e-mail in persistent storage;
means for processing the plurality of data nodes solely within non persistent storage, without requiring that information indicative of the e-mails be written to and then read from persistent storage during the processing of the e-mails, wherein said processing comprises, for each respective data node:
(i) determining a destination domain of the respective data node; and
(ii) adding the respective data node to a queue corresponding to the destination domain of the respective data node when the queue exists; and
(iii) creating a queue corresponding to the destination domain and adding the respective data node to the queue when the queue does not exist; and wherein
the means for processing further comprises:
selecting a first queue that contains data nodes;
retrieving e-mails corresponding to each of the data nodes in the first queue;
sending each of the e-mails corresponding to each of the data nodes in the first queue to a destination domain of the first queue; and
extinguishing the first queue.

85. (New) A computer system comprising:

means for obtaining a plurality of e-mails for processing;

means for forming a queue map comprising a plurality of queues, each queue in the plurality of queues associated with a specific domain, the queue map representing a plurality of destinations for the plurality of e-mails;

means for sending a plurality of e-mails to a specific destination in said plurality of destinations at a specific time; and

means for asynchronously looking up, during said sending, DNS information for a domain name using an asynchronous DNS resolver that operates from an offline DNS cache that is periodically updated, for a different destination in said plurality of destinations, to be sent at a future time.

86. (New) A computer system comprising:

means for obtaining a plurality of e-mails for processing;

means for forming organization information about said plurality of e-mails, wherein said organization information represents a plurality of queues, each queue in said plurality of queues comprising e-mails in said plurality of e-mails that are intended for distribution to a common destination; and

means for selecting a first queue in said plurality of queues to send e-mails, based on characteristics of the e-mails in the first queue and, during the selecting, asynchronously looking up DNS information for a domain name using an asynchronous DNS resolver that operates from an offline DNS cache that is periodically updated, for a second queue in said plurality of queues, different than the first queue.